AFRL-IF-RS-TM-2006-7
Final Technical Memorandum
November 2006

AUGMENTING SAT SOLVERS FOR NETWORK CONFIGURATION/PLANNING

Princeton University

Sponsored by
Defense Advanced Research Projects Agency
DARPA Order No. AOAB04 (00)

APPROVED FOR PUBLIC RELEASE; DISTRIBUTION UNLIMITED.

STINFO COPY

AIR FORCE RESEARCH LABORATORY
INFORMATION DIRECTORATE
ROME RESEARCH SITE
ROME, NEW YORK

NOTICE AND SIGNATURE PAGE

AFRL-IF-RS-TM-2006-7 HAS BEEN REVIEWED AND IS APPROVED FOR PUBLICATION IN ACCORDANCE WITH ASSIGNED DISTRIBUTION STATEMENT.

FOR THE DIRECTOR:

/s/	/s/
ROGER J. DZIEGIEL, JR. Work Unit Manager	JOSEPH CAMERA Chief, Information & Intelligence Exploitation Division Information Directorate

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

1. REPORT DATE *(DD-MM-YYYY)*	2. REPORT TYPE	3. DATES COVERED *(From - To)*
NOV 2006	Final	Apr 06 – Jun 06

4. TITLE AND SUBTITLE
AUGMENTING SAT SOLVERS FOR NETWORK CONFIGURATION/ PLANNING

5a. CONTRACT NUMBER FA8750-06-1-0164

5b. GRANT NUMBER

5c. PROGRAM ELEMENT NUMBER 62304E

6. AUTHOR(S)
Sharad Malik

5d. PROJECT NUMBER AB04

5e. TASK NUMBER 00

5f. WORK UNIT NUMBER 01

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)
Princeton University
1 Nassau Hall
Princeton New Jersey 08544-0001

8. PERFORMING ORGANIZATION REPORT NUMBER
N/A

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)

Defense Advanced Research Projects Agency	AFRL/IFED
3701 North Fairfax Drive	525 Brooks Road
Arlington Virginia 22203-1714	Rome New York 13441-4505

10. SPONSOR/MONITOR'S ACRONYM(S)

11. SPONSORING/MONITORING AGENCY REPORT NUMBER
AFRL-IF-RS-TM-2006-7

12. DISTRIBUTION AVAILABILITY STATEMENT
APPROVED FOR PUBLIC RELEASE; DISTRIBUTION UNLIMITED. PA #06-765

13. SUPPLEMENTARY NOTES

14. ABSTRACT
This project explored the possibility of alternate encodings of the planning problem and extensions to satisfiability (SAT) solvers that can better capture the constraints and objectives for network configurations. For example, the ability to directly deal with arithmetic constraints dealing with configuration cost, or probabilistic failure modes may lead to more compact encodings that are then dealt with more sophisticated decision procedures. The end goal is to capture the constraints and objectives of network planning in a decision problem and then solve these using efficient decision procedures in a scalable way.

15. SUBJECT TERMS
Satisfiability solvers, Quantified boolean formula, Modeling network configuration, Automatic network configuration and planning

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
a. REPORT	b. ABSTRACT	c. THIS PAGE			Roger J. Dziegiel, Jr.
U	U	U	UL	7	19b. TELEPONE NUMBER (*Include area code*)

Project: Augmenting SAT Solvers for Network Configuration/Planning

Principle Investigator: Sharad Malik, Princeton University

Overview: Network configuration problems need to determine network configurations, as well as network transformation plans that are robust in the face of adversarial action. Part of the overall research being conducted at Telcordia explores the possibility of modeling the objectives and constraints using a relational logic based modeling language (Alloy developed at MIT). These language descriptions are being currently compiled to propositional logic and their satisfiability checked using efficient contemporary SAT solvers (zchaff developed in our group at Princeton). However, this path currently does not seem scalable for large networks. This project will explore the possibility of alternate encodings of the planning problem and extensions to SAT solvers that can better capture the constraints and objectives. For example, the ability to directly deal with arithmetic constraints dealing with configuration cost, or probabilistic failure modes may lead to more compact encodings that are then dealt with more sophisticated decision procedures. The end goal is to capture the constraints and objectives of network planning in a decision problem and then solve these using efficient decision procedures in a scalable way.

Research Results: This work coupled with work done at Telcordia, MIT and Cornell as part of a collaborative project on “Automated Network Planning.” A final report of the collaborative effort has been presented to Chris Ramming of DARPA by Sanjai Narain of Telcordia on behalf of the group.

The following are the specific results for the work conducted at Princeton University by my group on this project.

- The modeling language being used to describe network configurations and constraints (Alloy from MIT) uses Boolean networks as its internal function representation. The reasoning process consists of converting the Boolean network into Conjunctive Normal Form (CNF) instances and calling a SAT solver to solve the CNF instance. We studied the benefits obtained from efficient encodings of the Boolean networks as CNF instances. The Boolean network can be converted to a circuit structure, e.g. the And/Inverter Graph (AIG). The AIG could be further simplified before being translated (encoded) into a CNF instance. Therefore the resulting CNF is expected to be smaller and easier for the SAT solver. Many AIG optimization techniques, e.g. structural hashing, SAT sweeping with observability don't cares, have been introduced in the past few years. We proposed a global AIG node resynthesis based approach that could further simplify the AIG dramatically. The algorithm details are available in the paper:
  Z. Fu, S. Malik, “AIG Rewriting using Global Node Resynthesis,” submitted for publication to the conference in Design, Automation and Test in Europe (DATE’06).

- In addition to the conventional encoding of simplified AIG into CNF instances, we can include the Circuit Observability Don’t Cares (Cir-ODC) information in the encoded CNF instance. The Cir-ODC information can help the SAT solver to dynamically prune irrelevant search spaces during the search and hence dramatically

improve the speed of the SAT solver. For some Alloy netconfig instances, we could see a large amount of the Cir-ODC information that could be captured during the CNF translation/encoding phase. Details of the Cir-ODC approach can be found in our previous paper:
Z. Fu, Y. Yu, S. Malik, “Considering Circuit Observability Don’t Cares in CNF Satisfiability,” in Design, Automation and Test in Europe (DATE’05).

- Several aspects of the netconfig instances need to consider extensions to SAT solvers. Some of there arise because of the need to consider numeric valued variables in addition to Boolean variables. We extended our SAT solver in several directions in response to that.

    - We investigated the extension of the SAT solver to handle the Minimum Cost Satisfiability (MinCostSAT) problem. A MinCostSAT instance is a classical SAT instance with a numeric cost of assigning each variable to be true or false. The objective is to find a satisfying assignment while minimizing the total cost incurred by the assignment. The MinCostSAT problem is essentially a generalization of many related problems, such as the Unate/Binate Covering problem and the Min/Max-Ones problems and it has been used in several other applications like FPGA routing, AI planning, etc. We proposed a SAT based branch-and-bound algorithm and implemented the MinCostChaff solver. The details of the solver can be found in the paper:
      Z. Fu, S. Malik, “Solving the MinCostSAT Problem using SAT Based Branch and Bound Search,” in International Conference on Computer-Aided Design (ICCAD’06), San Jose, CA.

      We have investigated the possibility of handling events using this MinCostSAT model and currently our solver is being used by the Alloy group at MIT.

    - Another extension of the SAT solver is to handle the Partial MAX-SAT (PM-SAT) problem. The PM-SAT problem sits in between the classical SAT and MAX-SAT problems. A PM-SAT instance consists of two types of clauses, the relaxable (or soft) clauses and the non-relaxable (hard) clauses. The objective is to find a variable assignment such that all non-relaxable clauses and maximum number of relaxable clauses are satisfied. This provides us with additional flexibility in modeling real world problems, particularly some AI planning and scheduling tasks in addition to netconfig instances where constraints could usually be prioritized. We proposed an iterative UNSAT Core elimination algorithm that guarantees the maximum number of relaxable clauses are satisfied. The detailed algorithm is described in the following paper:

Z. Fu, S. Malik, “On Solving The Partial MAX-SAT Problem,” in 9th International Symposium on the Theory and Applications of Satisfiability Testing (SAT’06), Seattle, WA.

  - o We also investigated the application of Satisfiability Modulo Theory (SMT) problems for network planning. An SMT problem is the satisfiability problem of propositional formulas on Boolean variables and predicates on background theories. An example of a background theory is linear arithmetic which results in linear arithmetic predicates on real and integer numbers like 3x + 4y > 5. An SMT problem can be conceived as an extension of a SAT problem. SMT problems can model real world problems that incorporate constraints with numeric values. Linear predicates are very useful in modeling deadlines of different tasks and their relationships, logistic constraints. These constraints have significant application in network planning. We proposed static and dynamic learning algorithms for SMT on linear predicates. Static learning is a pre-processing step that converts a part of complex linear predicates into simpler Boolean ones, which are much easier to reason with, while dynamic learning is incorporated into SMT solving to avoid repeatedly solving previously solved sub-problems. These algorithms improved the solving speed as well as the capacity of the current solvers. The details can be found in the paper:
    Y. Yu, Sharad Malik, “Lemma Learning in SMT on linear constraints, Proceedings,” International Conference on the Theory and Applications of Satisfiability Testing (SAT’06), Seattle, WA.

- o Besides these extensions to the SAT solvers, we also worked on fine-tuning the SAT solver to take advantage of certain special properties of the CNF instances translated from the network planning examples. One such property we found in our experiments is that such CNF instances usually contain much longer CNF clauses than normal CNF benchmarks. These long clauses have significant impact on several of the steps in the SAT solver, e.g. Boolean Constraint Propagation, clause deletion, etc. With some fine-tuning, the SAT solver could cut the execution time by half for some network planning instances.

- o The Cornell group worked on encodings of the network planning problem as adversarial games to be solved by a quantified Boolean solver (QBF). This prompted our work to study how a QBF solver can be improved for game instances when the initial specification is given as a circuit. In this case, we showed how it was possible to add don’t care information to the propositional part of the QBF instance (encoded in CNF) for significant speedup in the solution time for many instances. Details of this approach are in the following paper:
  D. Tang and S. Malik, “Solving Quantified Boolean Formulas with Circuit Observability Don’t Cares,” 9th International Conference on Theory and Applications of Satisfiability Testing (SAT), 2006

The above work conducted at Princeton was included with results from MIT, Cornell and Telcordia to study complete network planning/configuration problems. The final combined results were presented to Chris Ramming at DARPA.